

REMARKS

Claims 1, 2, 4-10, 12-17, 19, 20, 22 and 24-30 were previously pending in this application. No claims are amended, added or canceled. Claims 1, 2, 4-10, 12-17, 19, 20, 22 and 24-30 remain pending.

Correction to Serial Number

The Office correctly pointed out that the page header in the previous office action response included an incorrect serial number designation. Appropriate corrections have been made.

Objection to the Specification

The Abstract was objected to as comprising more than 150 words as prescribed by current regulations. The Abstract has been amended to fall within the 150 word limit.

35 U.S.C. § 103 Rejections

Claims 1, 2, 4, 5, 7, 8, 24-26 and 28-30

Claims 1, 2, 4, 5, 7, 8, 24-26 and 28-30 stand rejected under 35 U.S.C. 103(a) as being unpatentable over U.S. Patent No. 6,772,142 issued to Kelling et

al. (hereinafter "Kelling") in view of U.S. Patent No. 5,870,752 issued to Gibbons et al. (hereinafter "Gibbons"). Applicant respectfully traverses the rejection.

Claims 1, 2, 4, 5, 7 and 8

Claim 1 recites a method of maintaining a self-tuning histogram. The histogram is defined to have "a plurality of existing buckets arranged in a hierarchical manner and defined by at least two bucket boundaries that represent a range of attribute values, a bucket volume, and a bucket frequency that corresponds to a number of tuples having attribute values that fall in the bucket boundary range."

The method includes a step of "creating at least one new bucket in response to a query on the database, each new bucket having bucket boundaries corresponding to a range of tuple attribute values returned by the query and *a bucket frequency* corresponding to a number of tuples returned by the query."

The method further includes steps of "establishing a logical relationship between the new bucket and an existing bucket *such that the existing bucket is a parent bucket of the new bucket*" and "storing the self-tuning histogram that includes the new bucket in memory." (Emphasis added).

In the first phrase emphasized above, it is noted that a bucket has boundaries that correspond to a range of tuple attributes and a bucket frequency. The specification of the application (page 2, ¶ 1) discusses what a "bucket frequency" is as well as what it is not. "The frequency associated with a bucket, or bucket frequency, is an indication of the density of data within the bucket's boundaries, and should not be confused with the absolute value of the data within the bucket."

According to the claims, therefore, a bucket must be associated with an indication of the density of data within the bucket's boundaries.

Neither of the cited references associates such a bucket frequency with a bucket.

Kelling does not explicitly discuss bucket frequency. However, frequency may be inferred from the Kelling disclosure. Kelling discloses techniques for collecting and expressing geographically-referenced data. Kelling specifically discloses a system for identifying locations of sightings of particular species of birds to form a geographical map that identifies where such sightings have been made. A user may create a histogram from various locations shown on the map. For instance, if a user wanted to know how many red-headed woodpeckers were sighted in the state of New York over a period of time, the query could be

entered into the system and the system would provide the user with that information.

In such a system, frequency of a portion of a geographical area (a "bucket") would refer to absolute value of the data within the bucket. In other words, the system might answer the query by indicating there were one hundred sightings of the particular species in the particular geographical area. This is different from claim 1 and, therefore, Keller does not teach or suggest this element of claim 1.

Similarly, the term "frequency" as used in Gibbons relates to an absolute value of data within a bucket. In Gibbons, when one bucket includes a number of occurrences that exceeds a certain threshold, the bucket is divided into two separate buckets. The threshold indicates an absolute value of occurrences within the bucket. This is different from the requirements of claim 1, therefore, Gibbons does not teach or suggest this element of claim 1.

With regard to the second emphasized section from claim 1, above, when a new bucket is created, a parent-child relationship is established between the existing bucket and the new bucket, meaning that the new bucket is contained within the existing bucket (i.e. child buckets "and their bounding boxes are

disjoint and completely enclosed in" parent buckets (specification, page 10, ¶ 2 - page 11, ¶ 1)).

Neither of the cited references teaches or suggests establishing a parent-child relationship between buckets. In fact, establishing such a relationship is contrary to prior teachings in the art. Regarding Kelling, while one bucket may include other buckets (i.e. a "United States" bucket would include at least a "New York" bucket), the contents of the smaller, included buckets are not excluded from the larger buckets as they are in the present application. The Office admits that "Kelling does not disclose establishing a logical relationship between the new bucket and the existing bucket such that the existing bucket is a parent bucket of the new bucket." (Office Action, page 4, ¶ 1).

The Office states that Gibbons (Figs 4 and 5, col 10, lines 25-38) discloses this particular element. Applicant disagrees. These features of Gibbons show and discuss splitting one bucket to create two buckets if a threshold number of occurrences within a bucket is attained. Additionally, if the number of buckets exceeds a threshold, two buckets are merged to create a single bucket. The only relationship between buckets considered in the merge operation is whether two buckets are adjacent to each other. There is no

teaching or suggestion of treating buckets as having a parent-child relationship with other buckets.

Therefore, neither Kelling nor Gibbons teaches or suggests this particular element of claim 1.

Accordingly, claim 1 is not taught or suggested by either reference or a combination therefore. Therefore, the rejection of claim 1 should be withdrawn.

Claims 2, 4, 5, 7 and 8 depend from claim 1 and are allowable at least by virtue of that dependency. Therefore, the rejections of these claims should be withdrawn.

Claims 24-26 and 28-29

Claim 24 recites a histogram tuning system that comprises "a component that receives a histogram having at least a parent bucket" and "a tuning component that iteratively populates the parent bucket with a child bucket, as a function of query results, wherein the child bucket is completely contained within the parent bucket."

As previously discussed in the response to the rejection of claim 1, neither of the cited references nor a combination thereof teaches or suggests establishing a parent-child relationship between buckets, especially in light of the way in which Applicant has described a child bucket in the specification.

Therefore, claim 24 is allowable over the cited references and the rejection thereof should be withdrawn.

Claims 25-26 and 28-29 depend from claim 1 and are allowable at least by virtue of that dependency. Therefore, the rejections of these claims should be withdrawn.

Claim 30

Claim 30 recites a "database histogram tuning system" that comprises "means for receiving a bucket from a histogram" and "means for iteratively populating the bucket with a child bucket, as a function of query results, such that the child bucket is fully contained within the received bucket."

Claim 30 is similar to claim 24 and the analysis of the rejection with regard to the references is also similar. For the same reasons as discussed above, a combination of Kelling and Gibbons does not render claim 30 obvious.

Accordingly, claim 30 is allowable over the cited references and the rejection thereof should be withdrawn.

Claims 9, 10, 12-17, 19, 20 and 22

Claims 9, 10, 12-17, 19, 20 and 22 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over the combination U.S. Patent No. 6,507,840

issued to Ioannidis et al (hereafter Ioannidis) in view of Gibbons. Applicant respectfully traverses the rejection.

Claims 9, 10 and 12-15

Claim 9 recites a method a method of maintaining a self-tuning histogram. The histogram has "a plurality of existing parent buckets arranged in a hierarchical manner and defined by at least two bucket boundaries that represent a range of attribute values, a bucket volume, and a bucket frequency that corresponds to a number of tuples having attributes that fall in the bucket boundary range."

The method comprises "examining the results of a query executed on the database," "creating at least one candidate hole in the histogram based on the results of the query such that the candidate hole has boundaries corresponding to a range of attribute values returned by the query and a frequency corresponding to a number of tuples returned by the query," and "modifying the boundaries of each candidate hole such that the boundaries of the modified hole are completely contained within the boundaries of at least one existing parent bucket and do not partially intersect the boundaries of any existing bucket."

The method further includes "creating a new child bucket that has a child frequency in the histogram corresponding to each modified hole" and "storing the modified self-tuning histogram in one or more computer-readable media."

The Office Action admits that "Ioannidis ... does not disclose (b) creating at least one candidate hole in the histogram based on the results of the query such that the candidate hole has boundaries corresponding to a range of attribute values returned by the query and a frequency corresponding to a number of tuples returned by the query."

Applicant argues that Gibbons does not disclose the element either. As previously discussed, Gibbons does not use the term "frequency" in the same manner as described in the present specification. Also, when Gibbons discusses creating two buckets out of a single bucket, it only considers the number of occurrences included in the single bucket, which is different than what is required by claim 9.

In addition, claim 9 requires that the boundaries of each candidate hole are modified so that such boundaries are **"completely contained within the boundaries of at least one existing bucket...."** Applicant is unclear whether the Office cites Ioannidis or Gibbons for the element (see page 10, first full paragraph - [Gibbons, col 12, line 49, col 15, line 40-42] although Ioannidis

previously referenced at the beginning of the paragraph). Applicant disagrees and claims that neither references discloses this element of claim 9.

When Gibbons discloses creating a new bucket, it discusses creating two new buckets from a single existing bucket. Gibbons does not, however, disclose that one of the new buckets be completely contained within the boundaries of the existing bucket as required by claim 9.

Accordingly, claim 9 is allowable over the cited references and the rejection thereof should be withdrawn.

Claims 10 and 12–15 depend from claim 9 and are allowable at least by virtue of that dependency. Therefore, the rejections of these claims should be withdrawn.

Claims 16, 17 and 19

Claim 16 recites a bucket has having “at least two bucket boundaries that represent a range of attribute values, a bucket volume, and a bucket frequency that corresponds to a number of tuples having attribute values that fall in the bucket boundary range. . . .

As previously discussed, neither of the cited references utilize the term “frequency” as the present application.

Claim 16 also recites "creating at least one candidate hole in the histogram based on the results of the query such that the candidate hole has boundaries corresponding to a range of attribute values returned by the query and a frequency corresponding to a number of tuples returned by the query" and "modifying the boundaries of each candidate hole such that the boundaries of the modified hole are completely contained within the boundaries of at least one existing parent bucket and do not partially intersect the boundaries of any existing bucket."

As previously discussed, neither of the references nor a combination thereof teaches or suggests either or both of these elements. Here, creating a new bucket entails that the new bucket is disjoint from and completely included in another bucket. Nothing of this sort is disclosed in either of the cited references.

Accordingly, claim 17 is allowable over the cited references and the rejection should be withdrawn.

Claims 17 and 19 depend from claim 16 and are allowable at least by virtue of that dependency. Therefore, the rejections of these claims should be withdrawn.

Claim 20

Claim 20 is similar to claim 16 in that it recites "means for creating at least one candidate hole in the histogram based on the results of the query such that the candidate hole has boundaries corresponding to a range of attribute values returned by the query and a frequency corresponding to a number of tuples returned by the query." It also recites "means for modifying the boundaries of each candidate hole such that the boundaries of the modified hole are completely contained within the boundaries of at least one existing parent bucket and do not partially intersect the boundaries of any existing bucket."

For the same reasons as recited in the response to the rejection of claim 16, claim 20 is allowable over the cited references and the rejection of claim 20 should be withdrawn.

Claim 22

Claim 22 is similar to claim 16 and 22 and recites, inter alia, "a component that creates at least one candidate hole in the histogram based on the results of the query such that the candidate hole has boundaries corresponding to a range of attribute values returned by the query and a frequency corresponding to a number of tuples returned by the query." Claim 22 also discloses "a component that modifies the boundaries of each candidate

hole such that the boundaries of the modified hole are completely contained within the boundaries of at least one existing parent bucket and do not partially intersect the boundaries of any existing bucket."

For the same reasons as recited in the response to the rejection of claim 16, claim 22 is allowable over the cited references and the rejection of claim 22 should be withdrawn.

Claims 6 and 27

Claims 6 and 27 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over the combination Kelling and Gibbons and further in view of Ioannidis. Applicant respectfully traverses the rejection.

Claim 6

Claim 6 depends ultimately, from claim 1 and is allowable at least by virtue of that dependency. Therefore, the rejection of claim 6 should be withdrawn.

Claim 27

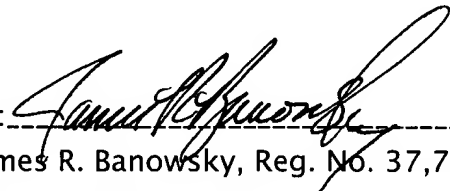
Claims 27 depends, ultimately, from claim 24 and is allowable at least by virtue of that dependency. Therefore, the rejection of claim 27 should be withdrawn.

CONCLUSION

In view of the foregoing amendments and remarks, this application should now be in condition for allowance. A notice to this effect is respectfully requested. If the Examiner believes, after this amendment, that the application is not in condition for allowance, the Examiner is encouraged to call the Applicant's attorney at the telephone number listed below.

If this response is not considered timely filed and if a request for an extension of time is otherwise absent, Applicant hereby requests any necessary extension of time. If there is a fee occasioned by this response, including an extension fee, please charge any deficiency to Deposit Account No. 50-0463.

Respectfully submitted,  
*Chaudhuri et al., Applicants*

By:   
James R. Banowsky, Reg. No. 37,773  
Microsoft Corporation  
One Microsoft Way, M/S 8/1446  
Redmond, Washington 98052-6399  
Telephone (425) 705-3539

Docket No. 163192.01  
Dated: 8/1/05